arXiv:1609.09525v1 [cs.DS] 29 Sep 2016

Multi-dimensional signal approximation with sparse structured priors using split Bregman iterations ☆

Y. Isaac[a,c], Q. Barthélemy[b], C. Gouy-Pailler[a], M. Sebag[c], J. Atif[d]

[a] *CEA, LIST, 91191 Gif-sur-Yvette Cedex, France*
[b] *Mensia Technologies S.A., 75015 Paris, France*
[c] *Université Paris-Sud, TAO, CNRS − INRIA − LRI, 91405 Orsay, France*
[d] *PSL, Université Paris-Dauphine, LAMSADE, CNRS, UMR 7243, 75775 Paris, France*

Abstract

This paper addresses the structurally-constrained sparse decomposition of multi-dimensional signals onto overcomplete families of vectors, called dictionaries. The contribution of the paper is threefold. Firstly, a generic spatio-temporal regularization term is designed and used together with the standard ℓ_1 regularization term to enforce a sparse decomposition preserving the spatio-temporal structure of the signal. Secondly, an optimization algorithm based on the split Bregman approach is proposed to handle the associated optimization problem, and its convergence is analyzed. Our well-founded approach yields same accuracy as the other algorithms at the state-of-the-art, with significant gains in terms of convergence speed. Thirdly, the empirical validation of the approach on artificial and real-world problems demonstrates the generality and effectiveness of the method. On artificial problems, the proposed regularization subsumes the Total Variation minimization and recovers the expected decomposition. On the real-world problem of electro-encephalography brainwave decomposition, the approach outperforms similar approaches in terms of P300 evoked potentials detection, using structured spatial priors to guide the decomposition.

Keywords: Structured sparsity, overcomplete representations, analysis prior, split Bregman, fused-LASSO, EEG denoising.

1. Introduction

In the last two decades, dictionary-based representations have been applied with success on a number of tasks, e.g. robust transmission with compressed

☆The work presented in this paper has been partially funded by DIGITEO under the grant 2011-053D.

Email addresses: isaac.yoann@gmail.com (Y. Isaac), q.barthelemy@gmail.com (Q. Barthélemy), cedric.gouy-pailler@cea.fr (C. Gouy-Pailler), michele.sebag@lri.fr (M. Sebag), jamal.atif@dauphine.fr (J. Atif)

sensing [14], image restoration [35], blind source separation [33] or classification [54] to name a few. Dictionary-based representations proceed by approximating a signal with a linear combination of elements, referred to as dictionary atoms, where the dictionary is either given based on the domain knowledge, or learned from a signal database [48].
The state-of-the-art mostly considers overcomplete dictionaries; in such cases the signal decomposition is not unique, requiring to select a decomposition with specific constraints or properties. One such property is the *sparsity*, where the signal decomposition involves but a few atoms; the trade-off between the number of such atoms and the approximation error is controlled via a weighted penalization term. Despite their good properties, sparse decompositions are sensitive w.r.t. the data noise [15], particularly so when the dictionary atoms are highly correlated.

This paper focuses on preserving the structure of the signal through the dictionary-based decomposition, thereby expectedly decreasing its sensitivity w.r.t. noise. The proposed approach is motivated by, and illustrated on, spatio-temporal multi-dimensional signals, referred to as multi-channel signals. The decomposition of multi-channel signals involves: i) decomposing each channel into the dictionary (see [41] for a survey); ii) ensuring that the structure of the multi-channel data is preserved in the multi-dimensional decomposition.

Following and extending previous work [27], this paper focuses on **structured dictionary-based decomposition**, where the dictionary-based representation preserves the structure of the signal, as follows. Assuming that each one of the mono-dimensional signals is structured (e.g. being recorded in consecutive time samples and being continuous w.r.t. time), the structured decomposition property aims at preserving the signal structure in the dictionary-based representation (e.g. requiring that the approximation of signals in consecutive time steps is expressed on same atoms with "close" weights). Formally, the structured decomposition property is enforced via considering a specific regularization term besides the data fitting term (minimization of the approximation error) and the standard ℓ_1 term (maximizing the decomposition sparsity).
The main contribution of this paper is to propose an efficient optimization scheme based on the split Bregman iterations for performing this structured decomposition of multi-dimensional signals. Implementation details including an efficient heuristic for hyper-parameters tuning is provided to accelerate the proposed algorithm. The proposed Multi-dimensional Sparse Structured Signal Approximation (Multi-SSSA) is assessed on synthetic signals for the fused-LASSO regularization obtained when the analysis term is a TV penalty. The approach is assessed in terms of i) its computational cost compared with the state-of-the-art smooth proximal gradient [9] approach; ii) its ability to recover the sparse structure of the initial signals compared with standard sparsity constraints and fused LASSO (ℓ_0, $\ell_{2,0}$, ℓ_1, $\ell_{2,1}$ and $\ell_1 + \ell_{2,1}$) even though the true structure of the signal is used to define the fused-LASSO regularization. Finally, Multi-SSSA is applied on electroencephalographic (EEG) signals to detect P300 evoked potentials. Using a data-driven prior, the structured decomposition approach is shown to effectively denoise the signal, resulting in a better

classification accuracy compared to the other regularizations.

The paper is organized as follows. Section 2 introduces the formal background. Related work is discussed in Section 3. Section 4 gives an overview of the proposed Multi-SSSA algorithm and the optimization strategy. The experimental validation of the approach on artificial data in terms of computational cost and structure recovery is presented and discussed respectively in Sections 5.2 and 5.3. Section 6 presents the denoising application on EEG signals. The paper concludes with a general discussion and points out some perspectives in Section 7.

Notations.. In the following, the j-th column of a matrix X is written $X(j)$, the i-th row $X^T(i)$, and the i-th element of j-th column $X(i,j)$. I_n stands for the identity matrix of size n. The ℓ_p matrix norm is defined as $\|X\|_p = (\sum_i \sum_j |X(i,j)|^p)^{\frac{1}{p}}$, with $p = 2$ corresponding to the classical Frobenius norm, and the $\ell_{p,q}$ mixed norm is defined as $\|X\|_{p,q} = (\sum_i (\sum_j |X(i,j)|^p)^{\frac{q}{p}})^{\frac{1}{q}}$.

2. Sparse structured decomposition problem

2.1. General problem

Let $Y = [Y(1), \ldots, Y(T)] \in \mathbb{R}^{C \times T}$ be a matrix of T ordered (e.g. corresponding to consecutive samples) C-dimensional signals, and $\Phi \in \mathbb{R}^{C \times N_\Phi}$ an overcomplete dictionary of N_Φ normalized atoms ($N_\Phi \gg C$). We consider the following linear model:

$$\begin{aligned} Y(t) &= \Phi X(t) + E(t), \quad t \in \{1, \ldots, T\} \ , \\ Y &= \Phi X + E \ , \end{aligned} \tag{1}$$

in which $X = [X(1), \cdots, X(T)] \in \mathbb{R}^{N_\Phi \times T}$ is the decomposition matrix and $E = [E(1), \cdots, E(T)] \in \mathbb{R}^{C \times T}$ stands for a Gaussian noise matrix.

The sparse structured approximation problem consists in decomposing each signal $Y(t), t \in \{1, \cdots, T\}$ onto the dictionary such that the joint decomposition X reflects the underlying structure of Y. This decomposition aims at removing the Gaussian noise present in these signals and at separating theirs components into target and non-target ones (depending on the considered application). The structured decomposition problem is formalized as the minimization of the objective function:

$$\min_{X \in \mathbb{R}^{N_\Phi \times T}} \|Y - \Phi X\|_2^2 + \lambda_1 \|X\|_1 + \lambda_2 \|XP\|_1 \ , \tag{2}$$

with λ_1, λ_2 the regularization coefficients and $P \in \mathbb{R}^{T \times N_P}$ a matrix encoding the prior knowledge about the signal structure. The use of the $\|XP\|_1$ regularization term can be interpreted in terms of sparse analysis [17]. Classically, the sparse analysis decomposition problem is formalized as follows:

$$\min_{X \in \mathbb{R}^{C \times T}} \|Y - X\|_2^2 + \lambda_1 \|\Psi X\|_1 \ , \tag{3}$$

whereas the synthesis problem writes as $\min_X \|Y - \Phi X\|_2^2 + \lambda_1 \|X\|_1$. In the analysis setting, atoms in the dictionary Ψ are viewed as filters, on which the projection of the decomposed signals are required to be sparse. Along this line, the P matrix can be viewed as a set of linear filters, on which the projections of the decomposed signals are required to be sparse, thereby preserving the regularities given from prior knowledge about the application domain. The matrix P can be learned from a set of signals [39, 42].

The present paper considers a synthesis formulation of the problem, regularized by an analysis term. In order to illustrate the interest of such decomposition, the particular case of the multi-dimensional fused-LASSO is described below.

2.2. Multi-dimensional fused-LASSO problem

In the particular case of a piecewise constant prior, a block-wise decomposition is expected. Formally this structure encoded by coefficients is defined as follows:

$$\forall n \in \{1, \cdots, N_\Phi\}, \quad X^T(n) = \sum_{m=1}^{M_n} \alpha_m^n \mathbf{1}_{\kappa_m^n} \ , \tag{4}$$

with $\{\kappa_m^n, \ \forall m \in \{1, \cdots, M_n\}\}$ a partition of $\{1, \cdots, T\}$ corresponding to the M_n blocks of the coefficients associated with the n-th atom of X, and $\{\alpha_m^n \in \mathbb{R}, \ \forall m \in \{1, \cdots, M_n\}\}$ the coefficients associated with the blocks.
In the studied model, adding an ℓ_1 and a multi-dimensional TV regularization terms allows to enforce such prior. The TV term leads to a sparse decomposition gradient, thus preserving signal singularities and data edges, and supporting the detection of abrupt changes [43]. The analysis term is then written: $\|XP^{\mathrm{TV}}\|_1 = \sum_{t=2}^{T} \|X(t) - X(t-1)\|_1$, where

$$P^{\mathrm{TV}} = \begin{pmatrix} -1 & & & \\ 1 & -1 & & \\ & 1 & \ddots & \\ & & \ddots & -1 \\ & & & 1 \end{pmatrix} \in \mathbb{R}^{T \times T-1} \ .$$

The Figure 1 depicts an example of time-series coefficients obtained in this case, *i.e.* with sparse variations, for a given atom of the dictionary (in grey). This combination of regularization terms is known as the fused-LASSO and has been chosen for the evaluation of the proposed algorithm on synthetic data because of its interest for various applications: frequency hopping [1], geophysical studies [19], multi-task learning [58], trend analysis [32], covariance estimation [12], analysis of association between genetic markers and traits [31] or change point detection [6], among others.

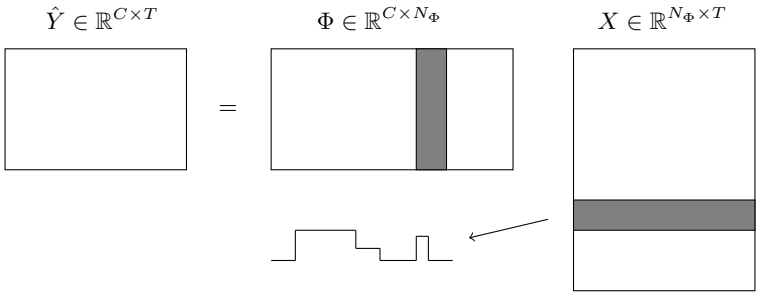

Figure 1: Decomposition of a multi-dimensional time-series signal $\hat{Y} = \Phi X$, with a TV regularization enforcing a block-wise structure of the coefficients X.

If the fused-LASSO is a well-known model, it is important to notice that the introduced Multi-SSSA is able to solve the general problem of Eq. (2) for any matrix P, with an analytic or a data-driven content.

3. Related work

Quite a few methods have been designed in the last decade to achieve the sparse approximation [11] of multi-dimensional signals while the mono-dimensional case has been intensively studied. Two main approaches have been considered: greedy methods trying to approximate the solution of the ℓ_0 regularized problem [50] and convex optimization solvers working on the ℓ_1 relaxed problem [49, 23, 41]. The concept of structured sparsity then emerged from the integration of prior knowledge encoded as regularizations into these approximation problems [26, 29].
The proposed Multi-SSSA approach falls in the latter category with a regularization combining a classical sparsity term and an analysis one. Some analysis regularizations have been extensively studied like the TV one which has been introduced in the ROF model [43, 13] for image denoising, nonetheless in the context of dictionary-based decomposition, their study has only begun recently [44, 51, 36]. Their interest for retrieving the underlying sparse structure of signals has been shown in [8]. To the best of our knowledge, the combination of these regularization terms had not been studied except for the particular case of the fused-LASSO introduced in [46].

Despite the convexity of the associated minimization problem, the two ℓ_1 non-differentiable terms make it difficult to solve (by classical gradient approaches). Various approaches have been developed for solving the fused-LASSO problem. In their seminal work, Tibshirani *et al.* [46] transformed the problem to a quadratic one and used standard optimization tools. While this approach proved computationally feasible for small-sized problems since it relies on increasing the dimension of the search space, it does not scale efficiently with problem size. Path algorithms have then been developed: Hoefling proposed in

[24] a method solving this problem in the particular case of the fused-LASSO signal approximator ($\Phi = I_C$) and Tibshirani *et al.* [47] designed a path method for the generalized LASSO problem. More recently, scalable approaches based on proximal sub-gradient methods [34], ADMM[1][53] and split Bregman iterations [56] have been successfully applied to the mono-dimensional generalized fused-LASSO.
Concerning the multi-dimensional fused-LASSO, an efficient method has been proposed recently in [9] for multi-task regression. A proximal method [37] is applied to a smooth approximation of the fused-LASSO regularization terms in this study and can be considered for different analysis regularizations. This approach is the most comparable with the present work.

To conclude this state-of-the-art, an illustration of the different regularizations is plotted in Figure 2: ℓ_2 (top left), ℓ_1 [45] (top right), $\ell_{2,1}$ [57] (middle left), $\ell_1 + \ell_{2,1}$ [22] (middle right), TV [43] (bottom right) and $\ell_1 + \text{TV}$ (bottom right).

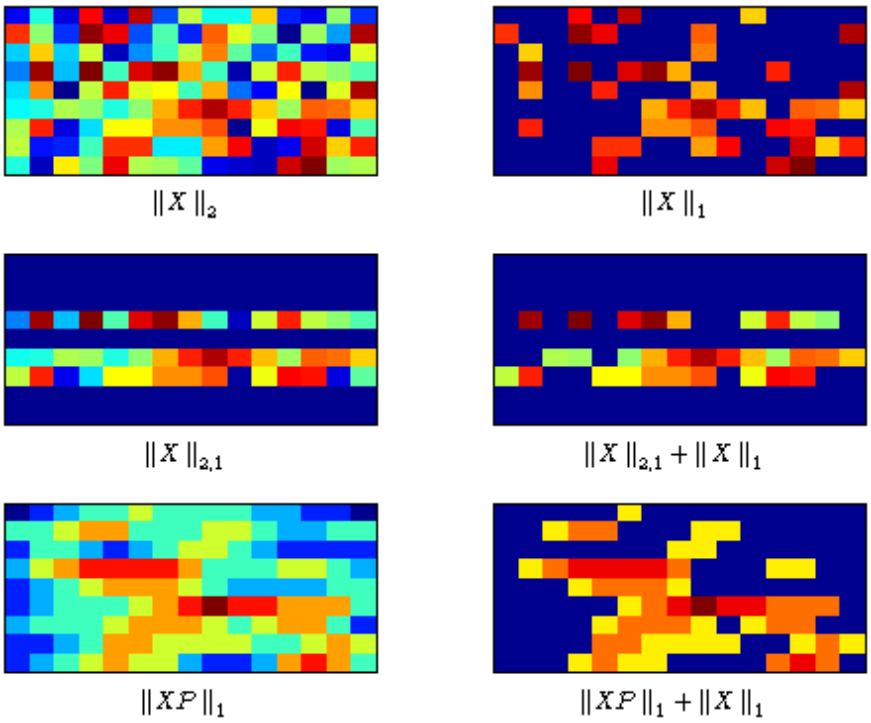

Figure 2: Illustration of the matrix $X \in \mathbb{R}^{N_\Phi \times T}$ with the different regularizations: ℓ_2 (top left), ℓ_1 (top right), $\ell_{2,1}$ (middle left), $\ell_1 + \ell_{2,1}$ (middle right), TV (bottom right) and $\ell_1 + \text{TV}$ (bottom right).

[1] Alternating Direction Method of Multipliers.

This figure also illustrates the interest of structurally constrained regularizations in various noisy situations. Many concrete applications entail extracting specific activities (signal space or target activities) from the original signals, therefore getting rid of various kind of noise, defined as spurious/non target activities. While target and non-target activities often share a few properties (statistical distribution, spatial and temporal structures), the introduction of structurally constrained regularizations aims at drastically limiting the resulting space to signals exhibiting the desired properties.

4. Optimization strategy

The split Bregman approach has been shown particularly well suited to ℓ_1 minimization problems [20] because of its ability to early detect zero and non-zero coefficients, which insures it a fast convergence on these problems. Thus, this optimization scheme has been chosen to solve the studied decomposition of Eq. (2).

4.1. Optimization Scheme

The above minimization problem is written as follows,

$$\min_{X\in\mathbb{R}^{N_\Phi\times T}} \|Y-\Phi X\|_2^2+\lambda_1\|X\|_1+\lambda_2\|XP\|_1 \ .$$

To setup the optimization scheme, let us first restate it as:

$$\begin{aligned}\min_{\substack{X\in\mathbb{R}^{N_\Phi\times T}\\ A\in\mathbb{R}^{N_\Phi\times T},B\in\mathbb{R}^{N_\Phi\times N_P}}} & \|Y-\Phi X\|_2^2+\lambda_1\|A\|_1+\lambda_2\|B\|_1\\ \text{s.t. } A=X \ \text{ and } \ B=XP \ . & \end{aligned} \tag{5}$$

This reformulation is a key step of the split Bregman method. It decouples the three terms and allows to optimize them separately within the iterations. To setup this iteration scheme, Eq. (5) is rewritten as an unconstrained problem:

$$\begin{aligned}\min_{\substack{X\in\mathbb{R}^{N_\Phi\times T}\\ A\in\mathbb{R}^{N_\Phi\times T},B\in\mathbb{R}^{N_\Phi\times N_P}}} & \|Y-\Phi X\|_2^2+\lambda_1\|A\|_1+\lambda_2\|B\|_1\\ & +\frac{\mu_1}{2}\|X-A\|_2^2+\frac{\mu_2}{2}\|XP-B\|_2^2 \ .\end{aligned}$$

Denoting i the current iteration, the split Bregman scheme [20] is then written:

$$\begin{aligned}(X^{i+1},A^{i+1},B^{i+1}) = & \operatorname*{argmin}_{\substack{X\in\mathbb{R}^{N_\Phi\times T}\\ A\in\mathbb{R}^{N_\Phi\times T},B\in\mathbb{R}^{N_\Phi\times N_P}}} \|Y-\Phi X\|_2^2+\lambda_1\|A\|_1+\lambda_2\|B\|_1\\ & +\frac{\mu_1}{2}\|X-A+D_A^i\|_2^2+\frac{\mu_2}{2}\|XP-B+D_B^i\|_2^2\\ D_A^{i+1} = & \ D_A^i+(X^{i+1}-A^{i+1})\\ D_B^{i+1} = & \ D_B^i+(X^{i+1}P-B^{i+1}) \ .\end{aligned} \tag{6}$$

This scheme is equivalent to the one obtained with the augmented Lagrangian method (or ADMM) when the constraints are linear [55]. Each iteration requires the solving of the primal problem before updating the dual variables.
Thanks to the split of the three terms, the minimization of the primal problem Eq. (6) can be performed iteratively by alternatively updating the following variables:

$$\begin{aligned} X^{i+1} = \underset{X\in\mathbb{R}^{N_\Phi\times T}}{\operatorname{argmin}} \quad & \|Y-\Phi X\|_2^2 + \frac{\mu_1}{2}\|X - A^i + D_A^i\|_2^2 \\ & + \frac{\mu_2}{2}\|XP - B^i + D_B^i\|_2^2 \end{aligned} \tag{7}$$

$$A^{i+1} = \underset{A\in\mathbb{R}^{N_\Phi\times T}}{\operatorname{argmin}} \quad \lambda_1\|A\|_1 + \frac{\mu_1}{2}\|X^{i+1} - A + D_A^i\|_2^2 \tag{8}$$

$$B^{i+1} = \underset{B\in\mathbb{R}^{N_\Phi\times N_P}}{\operatorname{argmin}} \quad \lambda_2\|B\|_1 + \frac{\mu_2}{2}\|X^{i+1}P - B + D_B^i\|_2^2 \ . \tag{9}$$

Empirically, it has been noted that only few iterations of this system are necessary for convergence [20]. In our implementation, this update is only performed once at each iteration of the global optimization algorithm.

Eq. (8) and Eq. (9) can be solved with the soft-thresholding operator [10]:

$$A^{i+1} = \underset{\frac{\lambda_1}{\mu_1}}{\text{SoftThreshold}}(X^{i+1} + D_A^i), \tag{10}$$

$$B^{i+1} = \underset{\frac{\lambda_2}{\mu_2}}{\text{SoftThreshold}}(X^{i+1}P + D_B^i) \tag{11}$$

with

$$(\underset{\lambda}{\text{SoftThreshold}}(X))(i,j) = \max\left(0,\ 1 - \frac{\lambda}{|X(i,j)|}\right) X(i,j).$$

Solving Eq. (7) requires the minimization of a convex differentiable function which can be performed via classical optimization methods. We propose here to solve it deterministically which is an original contribution of this work. Let us define H from Eq. (7) such as:

$$X^{i+1} = \underset{X\in\mathbb{R}^{N_\Phi\times T}}{\operatorname{argmin}} \ H(X) \ . \tag{12}$$

Differentiating this expression with respect to X yields:

$$\begin{aligned} \frac{d}{dX}H(X) = \quad & (2\Phi^T\Phi + \mu_1 I_{N_\Phi})X + X(\mu_2 PP^T) \\ & - 2\Phi^T Y + \mu_1(D_A^i - A^i) + \mu_2(D_B^i - B^i)P^T \ , \end{aligned} \tag{13}$$

The minimum $\hat{X} = X^{i+1}$ of Eq. (7) is obtained by solving $\frac{d}{dX}H(\hat{X}) = 0$, which is known as a Sylvester equation:

$$W\hat{X} + \hat{X}Z = M^i \ , \tag{14}$$

with $W = 2\Phi^T\Phi + \mu_1 I_{N_\Phi} \in \mathbb{R}^{N_\Phi \times N_\Phi}$, $Z = \mu_2 PP^T \in \mathbb{R}^{T\times T}$ and $M^i = 2\Phi^T Y + \mu_1(A^i - D_A^i) + \mu_2(B^i - D_B^i)P^T \in \mathbb{R}^{N_\Phi \times T}$.

In a general setting, solving efficiently a Sylvester equation can be time-consuming when the dimension of these matrices are large. A closed form solution derived from the vectorized formulation can be computed but it requires heavy calculations. One of the most used methods to deal with this issue, is the Bartels–Stewart algorithm [2] whose time complexity is $O(N_\Phi^3)$. In our particular case, the structures of the involved matrices ease this update.

Indeed, W and Z being real symmetric matrices, solving for $\hat{X}$ the Sylvester equation $W\hat{X} + \hat{X}Z = M^i$ is equivalent to solve for $\hat{X}'$ the following diagonal system

$$D_w\hat{X}' + \hat{X}'D_z = M^{i\prime} \ , \tag{15}$$

where

$$\begin{aligned} W &= FD_wF^T, & Z &= GD_zG^T \ , \\ \hat{X}' &= F^T\hat{X}G, & M^{i\prime} &= F^TM^iG \ . \end{aligned} \tag{16}$$

W and Z can be diagonalized in orthogonal bases (F and G) which allow to rewrite the Sylvester equation as follows:

$$FD_wF^T\hat{X} + \hat{X}GD_zG^T = M^i \ ,$$

and by applying the orthogonal properties of F and G the above expression is obtained.
The solution $\hat{X}'$ of the problem Eq.(15) is then computed as follows:

$$\forall t \in \{1, \cdots, T\} \quad \hat{X}'(t) = (D_w + D_z(t,t)I_{N_\Phi})^{-1}M^{i\prime}(t) \ ,$$

which is equivalent to

$$\begin{aligned} &\forall n \in \{1, \cdots, N_\Phi\}, \quad \forall t \in \{1, \cdots, T\} \ , \\ &(D_w(n,n) + D_z(t,t))\ \hat{X}'(n,t) = \ M^{i\prime}(n,t) \ , \end{aligned} \tag{17}$$

which can be computed with

$$\hat{X}' = M^{i\prime} \oslash O \ , \tag{18}$$

where $\oslash$ corresponds to an element-wise division and

$$O(n,t) = D_w(n,n) + D_z(t,t) \ . \tag{19}$$

$\hat{X}$ is then calculated as $\hat{X} = F\hat{X}'G^T$.

This last update can be numerically unstable when the elements of O are close to 0. As a consequence, μ_1 and μ_2 should be chosen carefully to avoid numerical instabilities. In addition, some terms can be precomputed to speed up the computations performed during each iteration. These implementation details and the choice of the penalty parameters are discussed in Section 4.3, and the full algorithm is summarized in Section 4.4.

4.2. Convergence

Thanks to the exact solving of the primal subproblem presented in the previous section, the convergence of the above scheme can be derived by drawing inspiration from the steps described in [7].

Theorem. *Assume that* $\lambda_1 \geq 0$, $\lambda_2 \geq 0$, $\mu_1 > 0$ *and* $\mu_2 > 0$, *then the following holds*

$$\begin{aligned} &\lim_{i\to\infty} \|Y - \Phi X^i\|_2^2 + \lambda_1 \|X^i\|_1 + \lambda_2 \|X^i P\|_1 \\ &\quad = \|Y - \Phi \hat{X}\|_2^2 + \lambda_1 \|\hat{X}\|_1 + \lambda_2 \|\hat{X} P\|_1 \end{aligned} \tag{20}$$

where $\hat{X}$ *denotes the solution of our problem.*
In addition, if our problem has a unique solution, from the convexity of $E(X) = \|Y - \Phi X\|_2^2 + \lambda_1 \|X\|_1 + \lambda_2 \|XP\|_1$ *and Eq. (20), we have*

$$\lim_{i\to\infty} X^i = \hat{X} \ . \tag{21}$$

Complete proof is available in Appendix A, demonstrating that split Bregman scheme converges to a solution of the convex problem. The uniqueness of the solution depend on the choice of P.

4.3. Implementation details and penalty parameters tuning

The terms W and Z in Eq. (14) are independent of the considered iteration i. Their diagonalizations can then be performed only once and for all as well as the computation of O. These diagonalizations are derived easily from those of $2\Phi^T\Phi$ and PP^T and can be realized off-line (and then pre-computed)

$$2\Phi^T\Phi = F\Delta_w F^T, \qquad PP^T = G\Delta_z G^T \tag{22}$$

$$D_w = \Delta_w + \mu_1 I_{N_\Phi} \text{ and} \qquad D_z = \mu_2 \Delta_z. \tag{23}$$

Thus, the update in Eq. (12) does not require heavy computations, even when the penalty parameters μ_1 and μ_2 change during the iterations (see below).

$Y_\Phi = 2F^T\Phi^T YG$ and $P_G = P^T G$ can also be pre-computed to avoid useless calculations. Besides, one can notice that the computational cost of each iteration depends on chain multiplications of three matrices (e.g $FX^{temp}G^T$). The computation time of these products depends on the order in which the multiplications are performed. The computational costs of the three chains appearing in the previously described scheme for both multiplication orders are presented in Figure 3.
The costs of the first two chains do not depend on the order of computation but the last one does. Hence, if $T \geq N_P$ then $F^T((D_B^i - B^{temp})P_G)$ is computed, otherwise $(F^T(D_B^i - B^{temp}))P_G$ is computed.

The choice of μ_1 and μ_2 has a crucial impact on the rate of convergence. On the one hand, they should be chosen to get a great conditioning of the primal problem. On the other hand, these parameters can be updated during the

Chain	Cost LR	Cost RL
$F^T(D_A^i - A^i)G$	$N_\Phi T(N_\Phi + T)$	$N_\Phi T(N_\Phi + T)$
FX^iG^T	$N_\Phi T(N_\Phi + T)$	$N_\Phi T(N_\Phi + T)$
$F^T(D_B^i - B^i)P_G$	$N_\Phi N_P(T + N_\Phi)$	$N_\Phi T(N_\Phi + N_P)$

Figure 3: Computation costs of multiplication chains. LR: ABC computed with $(AB)C$; RL: ABC computed with $A(BC)$.

iterations to improve this rate as it is classically done in augmented Lagrangian methods [4].

A poor conditioning of the primal update appears when these parameters are too small. The subproblem written in Eq. (7) becomes numerically unstable when the elements of the matrix $O(n,t) = D_w(n,n) + D_z(t,t)$ are close to 0. The eigenvalues of $W = 2\Phi^T\Phi + \mu_1 I_{N_\Phi}$ and $Z = \mu_2 PP^T$ are non-negative and their smallest values depend on the penalty parameters.
Let $\{\lambda_w(n), \forall n \in \{1, \cdots, N_\Phi\}\}$ and $\{\lambda_z(t), \forall t \in \{1, \cdots, T\}\}$ be respectively the eigenvalues of $2\Phi^T\Phi$ and PP^T, we have:

$$\min_{n,t} O(n,t) = \min_n \lambda_w(n) + \mu_1 + \mu_2 \min_t \lambda_z(t) \ . \tag{24}$$

When Φ is an overcomplete dictionary, $\min_n \lambda_w(n) = 0$ since this is the minimal singular value of Φ. Thus, depending on the eigenvalues of PP^T, μ_1 and μ_2 should be chosen carefully to avoid numerical instability.

An update of these parameters during the iterations can speed up the convergence. As seen before, the split Bregman scheme is equivalent to the augmented Lagrangian one when the constraints are linear. Thus, a common strategy used in the augmented Lagrangian scheme has been chosen here to update the penalty parameters: when the loss associated with a constraint does not decrease enough between two iterations, the corresponding parameter is increased. Formally, let $h_1(X,A) = \|X - A\|_2$ and $h_2(X,B) = \|XP - B\|_2$ be the constraints losses, the parameters are updated as follows,

$$\mu_1^i = \begin{cases} \mu_1^{i-1} & \text{if } \ h_1(X^i, A^i) < r_1 h_1(X^{i-1}, A^{i-1}) \\ \rho_1 \mu_1^{i-1} & \text{otherwise} \end{cases} \tag{25}$$

and

$$\mu_2^i = \begin{cases} \mu_2^{i-1} & \text{if } \ h_2(X^i, B^i) < r_2 h_2(X^{i-1}, B^{i-1}) \\ \rho_2 \mu_2^{i-1} & \text{otherwise} \end{cases} \tag{26}$$

where r_1, r_2 are the threshold parameters and ρ_1, ρ_2 are the ratios of the geometric progressions μ_1, μ_2.

To achieve a fast convergence, the initialization of these parameters should not enforce the constraints too strictly in the first iterations but these parameters must not be setup with a value too small in order to not affect the resolution of the primal problem. A heuristic is described here to ease the initialization of μ_1 and μ_2. This procedure has been shown to be empirically efficient:

1. define a search grid g for μ_1 and μ_2,
2. perform the first iteration of the optimization scheme for each couple of parameters $[g(j), g(l)]$,
3. evaluate $t_1(j,l) = \frac{\mu_1}{2} h_1(X(j,l)^1, A(j,l)^1)^2$ and $t_2(j,l) = \frac{\mu_2}{2} h_2(X(j,l)^1, B(j,l)^1)^2$ for each couple,
4. initialize μ_1^0 and μ_2^0 with:

$$\mu_1^0 = g(\operatorname*{argmax}_j \sum_l t_1(j,l)) \ ,$$

$$\mu_2^0 = g(\operatorname*{argmax}_l \sum_j t_2(j,l)) \ .$$

4.4. Full Multi-SSSA algorithm

The Multi-SSSA is summarized below as a pseudo-code procedure. To ease the understanding of the method, only the crucial steps are highlighted. The initialization and performance optimization steps just point out their corresponding equations. Note that only one subiteration is needed to solve the primal subproblem.

Parameters: λ_1, λ_2, μ_1^0, μ_2^0, ϵ, $iterMax$, $kMax$, r_1, r_2, ρ_1, ρ_2

procedure MULTI-SSSA(Y, Φ, P)
  Initialize D_A^0, D_B^0, X^0 and set $B^0 = X^0 P$, $A^0 = X^0$
  Diagonalize $2\Phi^T\Phi$ and PP^T to get Δ_w, Δ_z, F and G ▷ Eq. (22)
  Compute D_w, D_z, from Δ_w, Δ_z, μ_1^0 and μ_2^0) ▷ Eq. (23)
  Calculate O from D_w and D_z ▷ Eq. (19)
  Precompute $Y_\Phi = 2F^T\Phi^T Y G$ and $P_G = P^T G$
  $i = 0$
  while $i \leq iterMax$ and $\frac{\|X^i - X^{i-1}\|_2}{\|X^i\|_2} \geq \epsilon$ **do** ▷ stopping criteria
    $M' = Y_\Phi - F^T(\mu_1^i(D_A^i - A^i)G$
      $+ \ \mu_2^i(D_B^i - B^i)P_G)$
    $X^{i+1} = M' \oslash O$
    $X^{i+1} = F X^{i+1} G^T$ ▷ X primal update
    $A^{i+1} = \text{SoftThreshold}_{\frac{\lambda_1}{\mu_1^i}}(X^{i+1} + D_A^i)$ ▷ A primal update
    $B^{i+1} = \text{SoftThreshold}_{\frac{\lambda_2}{\mu_2^i}}(X^{i+1}P + D_B^i)$ ▷ B primal update
    $D_A^{i+1} = D_A^i + (X^{i+1} - A^{i+1})$ ▷ D_A dual update
    $D_B^{i+1} = D_B^i + (X^{i+1}P - B^{i+1})$ ▷ D_B dual update
    Compute μ_1^{i+1} and μ_2^{i+1} ▷ Eq. (25, 26)
    Update D_w, D_z and O ▷ Eq. (23, 19)
    $i = i + 1$
  end while
  return X^i
end procedure

5. Experiments on synthetic data

To evaluate the proposed method, two experiments have been performed on synthetic data. The computational time of the algorithm is first evaluated w.r.t. the state-of-the-art. Then, its ability to recover the underlying structure of signals is compared with other classical decomposition approaches. Both experiments have been carried out in the particular case of the fused-LASSO regularization (cf. Eq. (2) with $P = P^{\mathrm{TV}}$) on synthetic piecewise constant signals.

5.1. Data generation

Each of these piecewise constant signals $Y^k, \quad k \in \{1, \cdots, K\}$ has been synthesized from a built decomposition matrix X^k and a dictionary Φ with $Y^k = \Phi X^k$. The atoms of Φ have been drawn independently from a Gaussian distribution to create a random overcomplete dictionary, which consequently has a low coherence:

$$\max_{i,j \ \in\{1,\cdots,N_\Phi\},\ i\neq j} \Phi(i)^T \Phi(j) \approx 0.35 \ .$$

Each block-wise decomposition matrix X has been built as a linear combination of specific activities generated as follows:

$$\Theta_{n,t,d}(i,j) = \begin{cases} 0 & \text{if } i \neq n \\ \mathcal{H}(j - (t - \frac{d\times T}{2})) & \\ -\mathcal{H}(j - (t + \frac{d\times T}{2})) & \text{if } i = n \end{cases}$$

where $\Theta \in \mathbb{R}^{N_\Phi \times T}$, $\mathcal{H}$ is the Heaviside function, $n \in \{1, \cdots, N_\Phi\}$ the index of an atom, t the center of the activity and d its duration. Each decomposition matrix X could then be written:

$$X = \sum_{m=1}^{M} \alpha_m \Theta_{n_m,t_m,d_m} \ ,$$

where M is the number of activities appearing in one signal and the α_m stand for the activation weights. An example of generated signal is given in Figure 4.

5.2. Experimental assessment of computational time

The computational time of the Multi-SSSA algorithm is first evaluated. As noticed earlier, the most efficient method proposed for solving the multi-dimensional fused-LASSO is a smooth proximal-gradient method [9]. The TV analysis term is approached by a smooth penalty and an accelerated gradient descent is considered for minimizing the cost function. QP and SOCP formulations could also be considered but have been shown to be much slower than the proximal method [9]. Hence, the proposed optimization scheme is only compared here to this approach on the particular case of the multi-dimensional fused-LASSO.

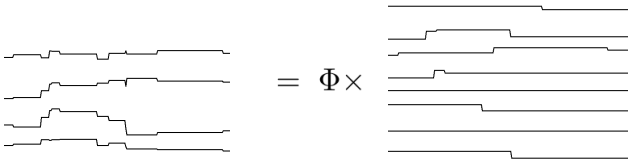

Figure 4: Example of synthesized signal $Y = \Phi X$, with $C = 4$ channels and $N_\Phi = 8$ atoms.

	C	N_Φ	T
T1	100	200	$50 \ldots_{50} 500$ $600 \ldots_{100} 1000$ $2000 \ldots_{1000} 6000$
T2	100	$50 \ldots_{50} 500$ $600 \ldots_{100} 1000$ $2000 \ldots_{1000} 5000$	300
T3	$50 \ldots_{50} 500$ $600 \ldots_{100} 1000$ $2000 \ldots_{1000} 8000$	200	300

Figure 5: Settings for speed comparison tests, where $m \ldots_n p = \{m+kn, \forall k \in \mathbb{N} \ s.t. \ m+kn \leq p\}$.

5.2.1. Experiments setup

Tests. To fairly compare these methods, the unique solution of our convex problem is calculated precisely beforehand. To compute this solution, we used the Multi-SSSA which is stopped when the relative change of the cost between two iterations is under 10^{-10}. Then, both methods are executed and stopped only when the relative difference between the current cost and the value of the loss at the optimum is under a fixed precision. Three experiments have been performed to assess the computational time of both methods when different dimensions of our problem are varying. Their respective setups are described in Figure 5. In addition, each test is realized for different values of the precision defined above: 10^{-4}, 10^{-5} and 10^{-6}.

Regularization parameters. The regularization parameter λ_1 has been fixed such that $\|Y - \Phi\hat{X}\|_2 / \|Y\|_2 \approx 0.1$ where $\hat{X}$ is the output of the optimization method and λ_2 determined by cross-validation on the distance between the decomposition matrices used to build the signals and those obtained as outputs.

Implementation details. Both methods have been implemented using MATLAB (64 bits). The experiments have been performed on a PC with 16GB RAM and a 8-core processor.
Concerning the proximal approach, the minimization is performed by the classical FISTA method as in [9]. More precisely, the variant named "FISTA with backtracking" in [3] has been chosen. The Lipschitzien coefficient L of the smooth term's gradient is approximated then with a variable following a geometrical progression: $L^i = \rho^k L^{i-1}$ with $\rho = 1.05$ and $L^0 = 1$. Besides, the parameter μ balancing the compromise between the analysis (TV) penalty and its smooth version is chosen such that the precision desired on the solution could be reached. A tight bound of the distance between the strict cost and the smooth one is known theoretically [9] and is proportional to μ. Consequently, to obtain $finalLoss \times precision \geq gapBound = K_{gap}\mu$ where $finalLoss$ is the minimum value of the cost function and $K_{gap} = \frac{1}{2}N_\Phi(T + N_P)$, μ is defined as follows:

$$\mu = 0.95 \times \frac{lossFinal \times precision}{K_{gap}} \ . \tag{27}$$

As for the introduced scheme, since the initial value of penalty parameters (obtained with the method presented in Section 4.3) have been observed to be stable (for the considered signals), they are computed off-line for each point of our tests on a 20×20 logarithmic grid. The update of these parameters is performed with $\rho_1 = \rho_2 = 1.05$ and $r_1 = r_2 = 0.95$. The diagonalizations of $2\Phi^T\Phi$ and PP^T are performed off-line.

5.2.2. Results and discussion

The results are presented in Figure 6. The execution times are displayed in a logarithmic scale. As expected, for both methods, the computation time is almost not affected by the number of channels. The observed increase of computation time is due to the computation of the stopping criteria at each iteration. Two more observations can be derived from these curves. Firstly, the split Bregman method is faster than the proximal gradient one in all cases presented here (for all precisions: 10^{-4}, 10^{-5} and 10^{-6}) and the curves have the same shapes. Secondly, the differences in speed between these approaches become more important when the desired precision becomes smaller. This last observation can be understood by noticing that the gradient's Lipschitzien coefficient of the smooth TV penalty presented in [9] is inversely proportional to the parameter μ. As this coefficient corresponds to the inverse of the FISTA gradient step coefficient, when μ becomes smaller (insuring a small gap) the gradient step becomes smaller and the method is then slower.

5.3. Experimental evaluation of sparse recovery on synthetic data

The interest of the studied regularization is now studied in a dictionary-based representation context. The performance of the described model to recover the underlying block-wise structures of artificial signals is assessed w.r.t. classical regularizations.

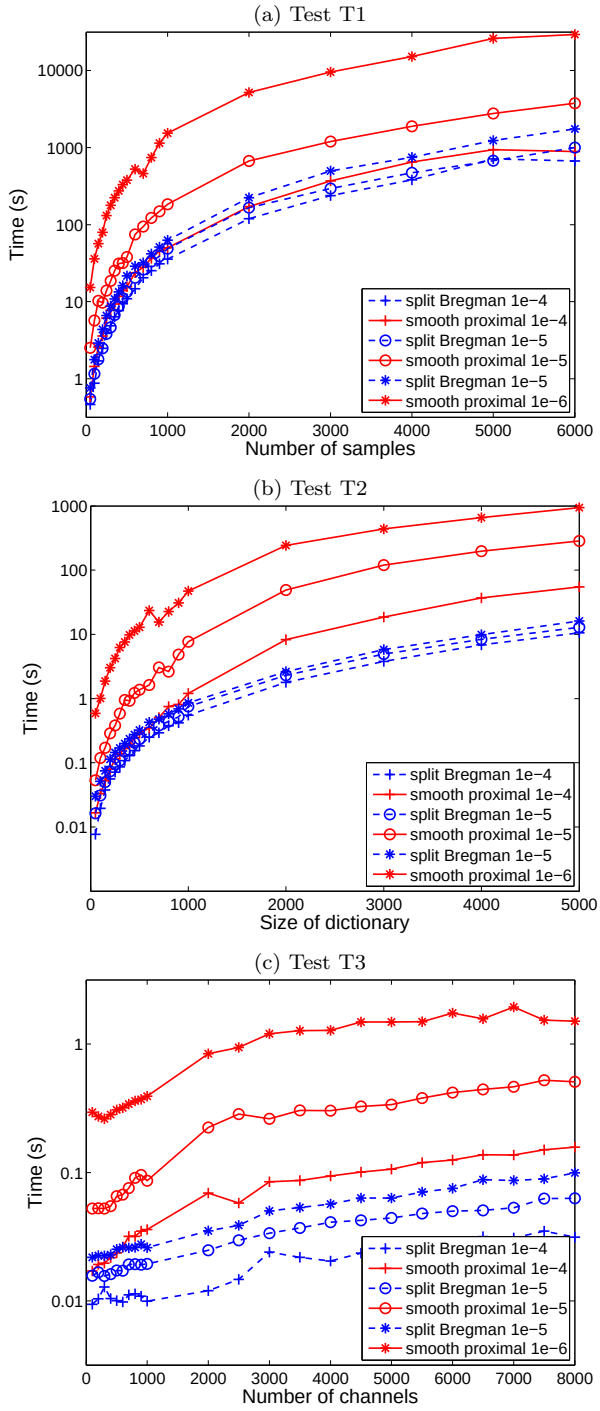

Figure 6: Speed comparison: split Bregman (blue) vs smooth proximal gradient (red), for different number of samples (a), size of dictionary (b) and number of channels (c).

5.3.1. Compared regularizations

Multi-SSSA is compared both with algorithms coding each signal separately with the ℓ_0 and ℓ_1 regularization terms and to methods performing the decomposition simultaneously with the $\ell_{2,0}$, $\ell_{2,1}$ and $\ell_{2,1}+\ell_1$ regularization terms. Regarding the ℓ_0 and $\ell_{2,0}$ constraints, the solutions are respectively given by the orthogonal matching pursuit (OMP) [38] and the simultaneous OMP (SOMP) [50]. The ℓ_1 solutions are obtained by the LARS method [16] and a proximal approach (FISTA [3]) has been chosen to deal with the $\ell_{2,1}$ and $\ell_1 + \ell_{2,1}$ regularizations [22]. The approximation problems with the ℓ_1 and $\ell_{2,1}$ regularization terms are respectively referred to as the LASSO [45] and group-LASSO [57] problems, where group-LASSO is used defining only one group by atom.

Among the compared algorithms, we have implemented the OMP and the SOMP, whereas the SPAMS[2] [30] toolbox has been used for the other methods.

5.3.2. Experimental settings

The goal of this experiment is to assess the ability of each model to retrieve the underlying structures of the designed piecewise constant signals. Since these performances vary with the number of activities M composing the signals and their duration d, this experience is performed for each point of the following grid of parameters:

- $M \in \{20, 30, \cdots, 110\}$,
- $d \sim \mathcal{U}(d_{min}, d_{max})$
  $(d_{min}, d_{max}) \in \{(0.05, 0.15), (0.15, 0.25), \cdots, (0.95, 1)\}$.

For each model and each point in the grid, the evaluation is carried out as follows:

- The set of built signals is split to create a training set allowing to determine the best regularization parameters and a test set designed to evaluate the performance with these parameters.
- For various regularization parameters, each signal Y of the training set is decomposed and the estimated decomposition matrix $\hat{X}$ is compared with the built one X using the following distance: $\varepsilon(X, \hat{X}) = \|X - \hat{X}\|_2 / \|X\|_2$. The parameters giving the best performances are chosen.
- Each signal of the test set is decomposed with the optimal parameters and the same distance is computed between the estimated decomposition matrix and the built one.

Other parameters of the signals construction are presented in Figure 7.

5.3.3. Results

For each point in the grid of parameters, the mean (among test signals) of the distance ε has been computed for each method and compared to the mean

[2] http://spams-devel.gforge.inria.fr/

Model		Activities	
$C = 20$	$T = 300$	$t \sim \mathcal{U}(0, T)$	$\alpha \sim \mathcal{N}(0, 2)$
$N_\Phi = 30$	$K = 100$	$n \sim \mathcal{U}(1, N_\Phi)$	

Figure 7: Parameters of the experiment of Section 5.3.

obtained by the Multi-SSSA algorithm. First, an analysis of variance shows that significant differences exist between the methods ($p \leq 0.05$). Paired t-tests (with Bonferroni corrections) have been performed to assess the significant differences between the couples of methods.
The results are displayed in Figure 8. In the ordinate axis, the number of patterns M increases from the top to the bottom and in the abscissa axis, the duration d grows from left to right. The top-left image displays the mean distances obtained by the Multi-SSSA algorithm, *i.e.* $\ell_1 + \text{TV}$. Unsurprisingly, the difficulty of finding the true decomposition increases with the number of patterns and their durations. The other figures present its performances compared to other methods by displaying the differences of mean distances in gray scale. These differences are calculated such that negative values (darker blocks) correspond to parameters where the introduced method outperforms the other one. The white diamonds correspond to non-significant differences of mean distances. Results of the OMP and the LASSO solver (ℓ_1) are very similar as well as those of the SOMP and the group-LASSO solver ($\ell_{2,1}$): they obtain the same pattern of performances on our grid of parameters. So, we only display here the matrices comparing the fused-LASSO regularization to those of the LASSO and group-LASSO models as well as with the regularization $\ell_1 + \ell_{2,1}$.

5.3.4. Discussions

First, concerning the comparison between the ℓ_1 (and ℓ_0) regularization terms and the $\ell_1 + \text{TV}$ one, it can be noted that similar results are obtained when only few atoms are active at the same time. It happens in our artificial signals when only few patterns have been used to create decomposition matrices and/or when the pattern durations are small. On the contrary, when many atoms are active simultaneously, the multi-dimensional fused-LASSO outperforms the LASSO model, allowing better retrieval of the block-wise structures of signals by using inter-signals prior information.
Concerning the $\ell_{2,1}$ (and $\ell_{2,0}$) regularization terms, results depend more on the duration of patterns. When patterns are longer, their performances are similar to the fused-LASSO one. On the contrary, when patterns have short/medium durations the group-LASSO model is outperformed. This is not surprising since these regularization terms select atoms for the entire duration of the signals.
As expected the $\ell_1 + \ell_{2,1}$ model combines the advantages of the ℓ_1 and $\ell_{2,1}$ regularization terms, having same performances as the fused-LASSO for both small number of patterns or long ones. In addition, its efficiency is better than the

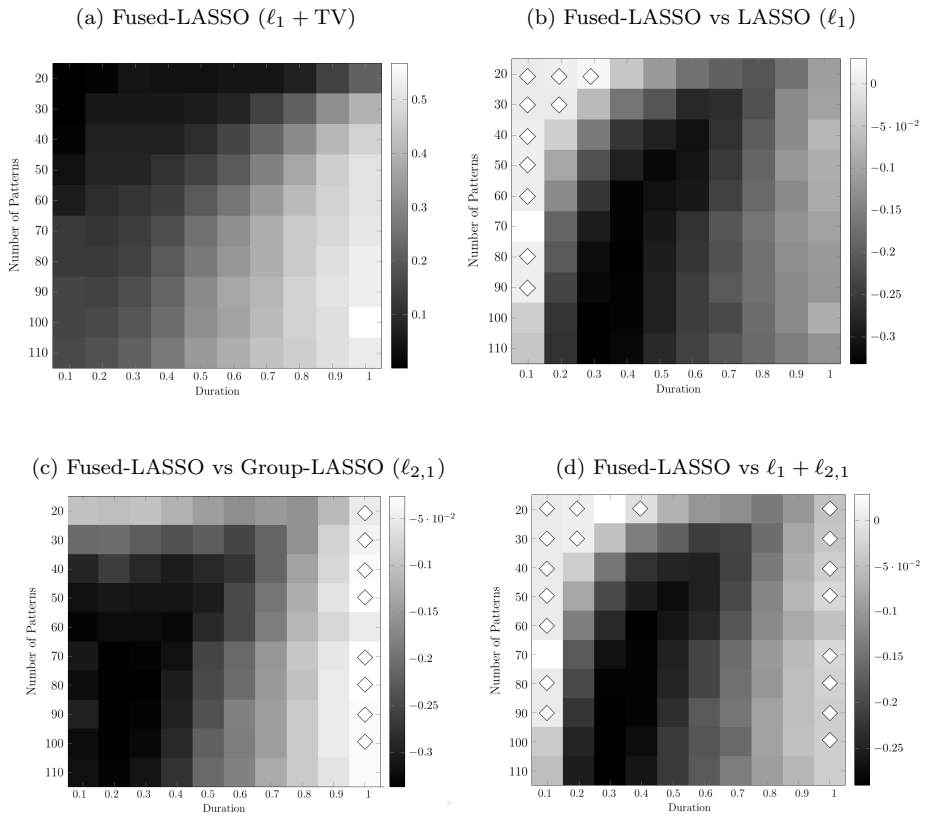
(a) Fused-LASSO (ℓ_1 + TV)
(b) Fused-LASSO vs LASSO (ℓ_1)
(c) Fused-LASSO vs Group-LASSO ($\ell_{2,1}$)
(d) Fused-LASSO vs $\ell_1 + \ell_{2,1}$
Number of Patterns
Duration
20 30 40 50 60 70 80 90 100 110
0.1 0.2 0.3 0.4 0.5 0.6 0.7 0.8 0.9 1
0.5 0.4 0.3 0.2 0.1
0 $-5 \cdot 10^{-2}$ −0.1 −0.15 −0.2 −0.25 −0.3

Figure 8: (a) Mean distances obtained with the Multi-SSSA. (b) Differences between the mean distances obtained with the Multi-SSSA and those obtained with the LASSO (ℓ_1) solver. (c) Differences between the mean distances obtained with the Multi-SSSA and those obtained with the Group-LASSO ($\ell_{2,1}$) solver. (d) Differences between the mean distances obtained with the Multi-SSSA and those obtained with the $\ell_1 + \ell_{2,1}$ solver. The white diamonds correspond to non-significant differences.

other studied regularizations in the middle of our grid even if it is still outperformed by the fused-LASSO which better detects the signals abrupt changes. Finally, this experiment illustrates the ability of the method to discriminate activities based on their statistical properties. Interferent signals, which share similar properties with a desired signal, will be split among distinct components if their spatial location of their temporal structure differ.

6. Application on EEG signals for P300 single-trial classification

The general model described in Eq. (2) of Section 2 can be applied to various contexts, with a non-analytic matrix P. This section is devoted to present the application of the studied regularization for the (unsupervised) denoising of real EEG data in a classification task: the detection of P300 evoked potentials.

6.1. Detection of P300 evoked potentials

The P300 is one of the most popular evoked potential used within Brain Computer Interfaces (BCI) systems. The P300 speller introduced in [18] allows for instance to spell words (letter by letter) by detecting such potentials after the presentation of visual stimuli. These potentials are usually elicited by presenting a rare target stimulus among common non-target ones (oddball approach). The P300 wave appears between 250 and 450 ms after the target stimulus and is mainly located in the parietal and the occipital lobes. Its latency and amplitude depend on various factors like the target-to-target intervals (see [40] for a review). The following experiment focuses on the single-trial detection of such brain activity.

Each EEG measurement is a spatio-temporal signal which can be studied as a sum of electrical activities emitted by different neural assemblies. These activities are characterized by specific time course and spatial patterns which depend on the locations of their sources and the orientations of the associated electrical activities.

In order to obtain plausible decompositions of such signals, the extracted components should respect the properties of these activities as the smoothness induced by the diffusion of electrical waves in the skull. The hypothesis in this experiment is that such plausible representations can be obtained via the studied regularization. The main components of these decompositions could then lead to a better identification of brain activities and in particular the components of the P300 potentials. Two types of noise must be removed: the sensors noise and the EEG background activities not related to the target activity. Since the components of the P300 potential are not known precisely, in order to evaluate the introduced regularization, the main components of the decompositions are used to reconstruct the signals. Then these denoised signals are used to identify P300 potentials.

6.2. Decomposition model

The previously studied model is considered with EEG signals $Y \in \mathbb{R}^{C\times T}$ recorded on T electrodes over C time samples[3]. These signals are decomposed thanks to the Multi-SSSA algorithm on a Gabor time-frequency dictionary $\Phi \in \mathbb{R}^{C\times N_\Phi}$ which is able to efficiently represent both transitory and oscilatory EEG components [52]. The analysis regularization matrix $P \in \mathbb{R}^{T\times N_P}$ is chosen as the dual of a spatial dictionary $\Phi_s \in \mathbb{R}^{N_P\times T}$ composed of N_P atoms. This dictionary is composed of realistic EEG topographies, horizontally concatenated since the regularization is applied with the dual on the lines of X (spatial dimension). An EEG topography corresponds to the values of the electrical potential of all electrodes at a given instant. The dual of this dictionary is approximated by the Moore–Penrose pseudo-inverse of Φ_s [17]. The following steps have been carried out to construct this spatial dictionary:

- a realistic head model has been built from MRI data and has been divided into voxels,
- for each voxel and different orientations of voxels' electrical activities, the associated EEG topographies have been computed by solving the EEG direct problem (with the software OpenMEEG [21]) and regrouped in a large spatial dictionary ($\approx$ 5000 elements),
- a subset of this large dictionary has been finally selected with a greedy approach such that its coherence does not exceed 0.9 ($N_P = 350$ elements).

Since the retrieval of the sources of these components is not the objective here, this latter step has been carried out to improve the conditioning of the decompositions and preserve a reasonable computational time (decomposition time of a signal around 2 seconds). Some topographies of Φ_s are presented in Figure 9.

6.3. Experimental settings

In this experiment, the studied regularization (denoted $MSSSA$) is compared to three others: ℓ_1, $\ell_{2,1}$ and $\ell_{2,1} + \ell_1$ to evaluate their ability to denoise P300 signals and ease their detection.

P300 data. The dataset IIb of the BCI Competition II [5] has been chosen to evaluate our model. These signals have been recorded on $T = 64$ electrodes at 240 Hz in several sessions. Trials have been extracted between 150 and 450 ms after each stimulus and band-filtered between 0.1 and 20 Hz with a fourth order Butterworth filter. The dataset is composed of three sessions and only the two first sessions have been used in this experiment.

[3]Model in Eq. (2) applies a structured prior along the time dimension whereas, in this section, the prior is applied along the spatial dimension. Thus, notations for spatial and temporal dimensions are inversed.

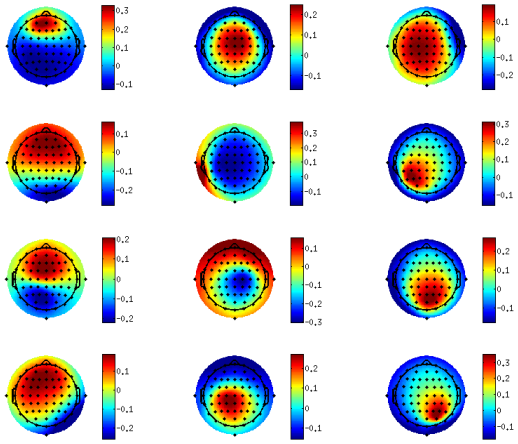

Figure 9: Visualization of several topographies of the spatial dictionary Φ_s.

Protocol. To evaluate the efficiency of each regularization, each trial Y^k is firstly decomposed on the dictionary Φ to obtain its decomposition matrix $\hat{X}^k$, then, a classification algorithm is applied to the set of the reconstructed signals $\{\hat{Y}^k = \Phi\hat{X}^k, k \in \{1, \ldots, K\}\}$. The BLDA (Bayesian Linear Discriminant Analysis) algorithm [25] has been chosen to perform the classification step. It has been shown particularly efficient for the detection of P300 evoked potentials.
The optimal regularization parameters of each model are learned on the second session of the dataset with a n-fold cross-validation, and the evaluation is then assessed on the first session of the dataset. This validation is performed for several values of n: 2 (500 signals in each fold), 5 (200 signals in each fold) and 10 (100 signals in each fold) in order to evaluate the regularizations efficiency for various sizes of the training set. Each cross-validation is performed 25 times, with folds selected randomly.
For this experiment, the $MSSSA$ has been stopped with the following criterion: $\|X_i - X_{i-1}\|_2 / \|X_i\|_2 \leq eps$ (with $eps = 10^{-6}$).

6.4. Results and discussion

The results of this experiment are presented in Figure 10. As expected, classification scores rise w.r.t. the number of folds and thus the size of the training set. Concerning the comparison between the regularizations, paired Wilcoxon signed-rank test have been performed between results obtained for each term, given the following results:

- the ℓ_1 regularization does not improve the scores obtained on raw signals ($p \geq 0.05$),
- $\ell_{2,1}$ and $\ell_{2,1} + \ell_1$ regularizations got the same results and improve significantly ($p \leq 0.005$) the scores obtained on raw signals.

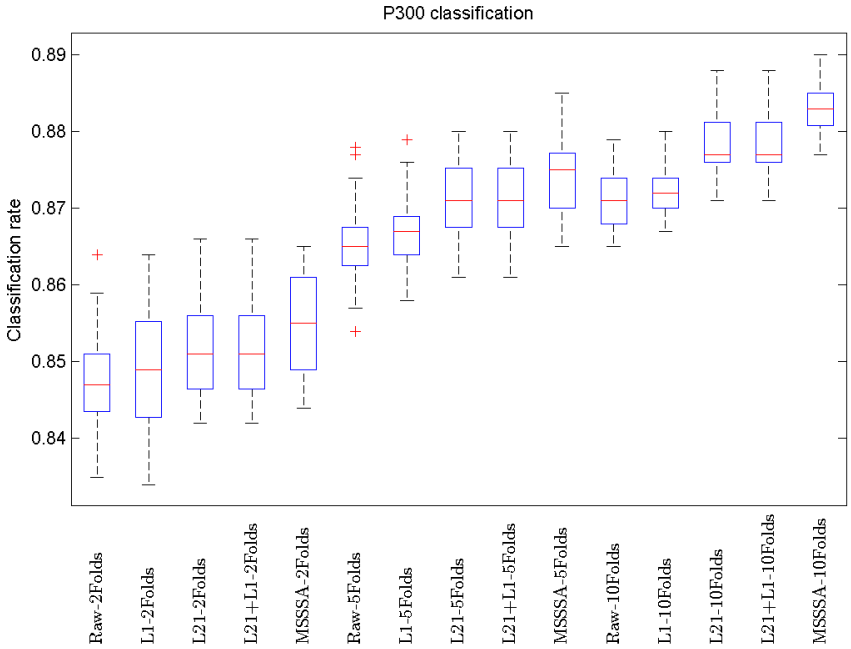

Figure 10: Comparison of classification rates between raw signals and rebuilt signals after decomposition with the $\ell_1(L1)$, $\ell_{2,1}(L21)$ and $\ell_{2,1}+\ell_1(L21+L1)$ and Multi-SSSA regularizations (referred to as $MSSSA$), for various number of folds in cross-validation.

- the introduced $MSSSA$ improve significantly ($p \leq 0.005$) the scores obtained with the $\ell_{2,1}$ regularization.

Only the regularizations enforcing a spatial structure in the decomposition coefficients ($\ell_{2,1}$ and $MSSSA$) allow here to improve the classification results by extracting plausible EEG components. The $\ell_{2,1}$ constraints all the channels to be decomposed on the same atoms, enforcing the choice of time-frequency atoms allowing to represent the time courses of all channels. The efficiency of this terms could be explained by noting that brain electrical activities are diffused by the skull and then affect a majority of the sensors. The proposed regularisation improves even more these classification scores by considering a data-driven spatial priors encoded in the analysis matrix P which guide the decomposition to more plausible P300 components.
Removing the noise $E^k = Y^k - \hat{Y}^k$, this decomposition can then be considered as a denoising step before P300 signals classification and more generally before EEG signals classification. In addition, the results could probably be improved by building the spatial dictionary from a head model of the subject on which the EEG are measured.

7. Conclusion and perspectives

The studied approach enforces structural properties of the signal decomposition on the dictionary through the regularization term $\|XP\|_1$. The overall optimization problem is efficiently and scalably handled through the split Bregman method.

Regarding the structural regularization term, a proof of principle of the approach is obtained in the particular case of the TV norm, conducive to the discovery of block-wise structures in the decomposition on overcomplete dictionaries: as shown in Section 5.3, the recovery of the block-wise structure of the input signals is significantly improved when using the ℓ_1 + TV regularization compared to other regularizations. In addition, the EEG experiment (Section 6) illustrates how the analysis term can be used to enforce structural properties on the decomposition in a data-driven way. It allows in this last case to denoise EEG signals and improve the single-trial detection of P300. The decompositions are guided to plausible components thank to an analysis matrix constructed from a dictionary of realistic EEG topographies.

The second contribution of the paper is the original split-Bregman method handling the underlying optimization problem, and its scalability compared with the state-of-the-art. As shown in Section 5.2, Multi-SSSA outperforms the smooth proximal gradient in terms of speed when the problem dimensions increase, and it is less sensitive w.r.t. the wanted precision. Furthermore, an empirically efficient heuristic procedure is proposed to adjust the penalty hyperparameters and thus preserve the efficiency of the algorithm.
This efficiency is explained from the split-Bregman ability to early detect zero coefficients (see [20], Appendix), thereby easily accommodating ℓ_1 regularization. The main scalability limitation of the proposed scheme comes from the diagonalization of the matrix in Eq. (22), with cubic complexity in its size; it therefore requires the dictionary size N_Φ and/or T to remain in the hundreds. A second limitation regards the memory complexity, and the storage of variables involved in the method.
To overcome these limitations, a method solving approximatively the primal subproblem Eq. (7) can be considered. Even if the exact solving of this subproblem allows to ensure the convergence of the proposed scheme, the split Bregman iterations have been shown empirically to converge even when the primal problem is not solved exactly. When N_Φ or N_P become really large, this option can significantly reduce the computation time.

Further perspectives and on-going work are primarily concerned with carrying out a full evaluation of the presented EEG application, in particular to examine the influence of the dictionaries choices (spatial and time-frequency) on the denoising efficiency. In addition, an extension of the microstates EEG model will be studied with a TV regularization matrix [28]. Another direction of research is concerned with approximating the solution of the subproblem in Eq. (12), instead of solving it exactly. A mid-term research perspective is to learn both the dictionary Φ and the regularities P from the data.
While not rigorously shown in this paper, we believe that the proposed frame-

work provides with a rich panel of possibilities to filter signal of interest from noise and interference in various context and applications.

Appendix A. Optimization scheme convergence

The convergence theorem (Section 4.2) of the introduced optimization scheme is here demonstrated by applying the convergence analysis of Osher *et al.* [7]. This theorem hold for $\lambda_1 \geq 0$, $\lambda_2 \geq 0$, $\mu_1 > 0$ and $\mu_2 > 0$.

The studied iterative algorithm consider at each iteration three convex sub-problems Eq. (7), (8), (9). The first order optimality condition of these problems gives:

$$\begin{aligned}
0 &= (2\Phi^T\Phi + \mu_1 I)X^{i+1} + \mu_2 X^{i+1}PP^T \\
&\quad - 2\Phi^T Y + \mu_1(D_A^i - A^i) + \mu_2(D_B^i - B^i)P^T \ , \\
0 &= \lambda_1 Q_A^{i+1} - \mu_1(D_A^i - A^{i+1} + X^{i+1}) \ , \\
0 &= \lambda_2 Q_B^{i+1} - \mu_2(D_B^i - B^{i+1} + X^{i+1}P) \ , \\
D_A^{i+1} &= D_A^i + (X^{i+1} - A^{i+1}) \ , \\
D_B^{i+1} &= D_B^i + (X^{i+1}P - B^{i+1}) \ ,
\end{aligned} \tag{A.1}$$

where $Q_A^{i+1} \in \partial\|A^{i+1}\|_1$ et $Q_B^{i+1} \in \partial\|B^{i+1}\|_1$.

In addition, the convexity of the main problem Eq. (5) insures the existence of an unique solution which respect the KKT conditions. The Lagrangian L of the problem could thus be written:

$$L = \|Y - \Phi X\|_F^2 + \lambda_1\|X\|_1 + \lambda_2\|XP\|_1,$$

and $\exists \hat{X}$ such that

$$\begin{aligned}
0 &= -2\Phi^T(Y - \Phi\hat{X}) + \lambda_1\hat{Q}_A + \lambda_2\hat{Q}_B P^T \ , \\
\hat{A} &= \hat{X}, \quad \text{and} \quad \hat{B} = \hat{X}P \ ,
\end{aligned} \tag{A.2}$$

where $\hat{Q}_A \in \partial\|\hat{A}\|_1$ et $\hat{Q}_B \in \partial\|\hat{B}\|_1$.

This solution is a fixed point of the optimization scheme and verifies:

$$\begin{aligned}
0 &= (2\Phi^T\Phi + \mu_1 I)\hat{X} + \mu_2\hat{X}PP^T - 2\Phi^T Y + \mu_1(\hat{D}_A - \hat{A}) + \mu_2(\hat{D}_B - \hat{B})P^T \ , \\
0 &= \lambda_1\hat{Q}_A - \mu_1(\hat{D}_A - \hat{A} + \hat{X}) \ , \\
0 &= \lambda_2\hat{Q}_B - \mu_2(\hat{D}_B - \hat{B} + \hat{X}P) \ , \\
\hat{D}_A &= \hat{D}_A + (\hat{X} - \hat{A}) \ , \\
\hat{D}_B &= \hat{D}_B + (\hat{X}P - \hat{B}) \ .
\end{aligned} \tag{A.3}$$

Substracting Eq. (A.3) from Eq. (A.1) we obtained the same system with

the errors variables:

$$\begin{aligned}\tilde{X}^i = X^i - \hat{X},\ \tilde{A}^i = A^i - \hat{A},\ \tilde{B}^i = B^i - \hat{B}\ ,\\ \tilde{D}_A^i = D_A^i - \hat{D}_B,\ \tilde{D}_B^i = D_B^i - \hat{D}_B\\ \tilde{Q}_A^i = Q_A^i - \hat{Q}_A,\ \tilde{Q}_B^i = Q_B^i - \hat{Q}_B\ .\end{aligned}$$

Performing the scalar product of the first line by $\tilde{X}^{i+1}$, the scalar product of the second line by $\tilde{A}^{i+1}$, the scalar product of the third line by $\tilde{B}^{i+1}$ and taking the square Frobenius norm of the last lines, we obtained the following system:

$$\begin{aligned}0 &= 2\|\Phi\tilde{X}^{i+1}\|_F^2 + \mu_1\|\tilde{X}^{i+1}\|_F^2 + \mu_2\langle\tilde{X}^{i+1}, \tilde{X}^{i+1}PP^T\rangle\\ &\quad + \mu_1(\langle\tilde{X}^{i+1}, \tilde{D}_A^i\rangle - \langle\tilde{X}^{i+1}, \tilde{A}^i\rangle) + \mu_2(\langle\tilde{X}^{i+1}, \tilde{D}_B^iP^T\rangle - \langle\tilde{X}^{i+1}, \tilde{B}^iP^T\rangle)\ ,\\ 0 &= \lambda_1\langle\tilde{A}^{i+1}\tilde{Q}_A^{i+1}\rangle - \mu_1(\langle\tilde{A}^{i+1}\tilde{D}_A^i\rangle - \|\tilde{A}^{i+1}\|_F^2 + \langle\tilde{A}^{i+1}, \tilde{X}^{i+1}\rangle)\ ,\\ 0 &= \lambda_2\langle\tilde{B}^{i+1}, \tilde{Q}_B^{i+1}\rangle - \mu_2(\langle\tilde{B}^{i+1}, \tilde{D}_B^i\rangle - \|\tilde{B}^{i+1}\|_F^2 + \langle\tilde{B}^{i+1}, \tilde{X}^{i+1}P\rangle)\ ,\\ \|\tilde{D}_A^{i+1}\|_F^2 &= \|\tilde{D}_A^i\|_F^2 + (\|\tilde{X}^{i+1}\|_F^2 + \|\tilde{A}^{i+1}\|_F^2\\ &\quad - 2\langle\tilde{X}^{i+1}, \tilde{A}^{i+1}\rangle) - 2\langle\tilde{D}_A^i, \tilde{X}^{i+1} - \tilde{A}^{i+1}\rangle\ ,\\ \|\tilde{D}_B^{i+1}\|_F^2 &= \|\tilde{D}_B^i\|_F^2 + (\|\tilde{X}^{i+1}P\|_F^2 + \|\tilde{B}^{i+1}\|_F^2\\ &\quad - 2\langle\tilde{X}^{i+1}P, \tilde{B}^{i+1}\rangle) - 2\langle\tilde{D}_B^i, \tilde{X}^{i+1}P - \tilde{B}^{i+1}\rangle\ .\end{aligned}$$

Summing the 3 first equations and slightly modifying the others, gives:

$$\begin{aligned}0 = &\ 2\|\Phi\tilde{X}^{i+1}\|_F^2 + \mu_1\|\tilde{X}^{i+1}\|_F^2 + \mu_2\langle\tilde{X}^{i+1}, \tilde{X}^{i+1}PP^T\rangle\\ &+ \lambda_1\langle\tilde{A}^{i+1}, \tilde{Q}_A^{i+1}\rangle + \lambda_2\langle\tilde{B}^{i+1}, \tilde{Q}_B^{i+1}\rangle\\ &+ \mu_1(\langle\tilde{X}^{i+1}, \tilde{D}_A^i\rangle - \langle\tilde{X}^{i+1}, \tilde{A}^i\rangle - \langle\tilde{A}^{i+1}, \tilde{D}_A^i\rangle + \|\tilde{A}^{i+1}\|_F^2 - \langle\tilde{A}^{i+1}, \tilde{X}^{i+1}\rangle)\\ &+ \mu_2(\langle\tilde{X}^{i+1}, \tilde{D}_B^iP^T\rangle - \langle\tilde{X}^{i+1}, \tilde{B}^iP^T\rangle - \langle\tilde{B}^{i+1}, \tilde{D}_B^i\rangle + \|\tilde{B}^{i+1}\|_F^2 - \langle\tilde{B}^{i+1}, \tilde{X}^{i+1}P\rangle)\ ,\\ \langle\tilde{D}_A^i, \tilde{X}^{i+1} - \tilde{A}^{i+1}\rangle = &\ \frac{1}{2}(\|\tilde{D}_A^{i+1}\|_F^2 - \|\tilde{D}_A^i\|_F^2 - \|\tilde{X}^{i+1} - \tilde{A}^{i+1}\|_F^2)\ ,\\ \langle\tilde{D}_B^i, \tilde{X}^{i+1}P - \tilde{B}^{i+1}\rangle = &\ \frac{1}{2}(\|\tilde{D}_B^{i+1}\|_F^2 - \|\tilde{D}_B^i\|_F^2 - \|\tilde{X}^{i+1}P - \tilde{B}^{i+1}\|_F^2)\ .\end{aligned}$$

and combining these equations and summing between $i = 1$ and $i = I$:

$$\begin{aligned}&\frac{\mu_1}{2}(\|\tilde{D}_A^1\|_F^2 - \|\tilde{D}_A^S\|_F^2) + \frac{\mu_2}{2}(\|\tilde{D}_B^1\|_F^2 - \|\tilde{D}_B^S\|_F^2)\\ &\quad = 2\sum_{i=1}^{S}\|\Phi\tilde{X}^i\|_F^2 + \sum_{i=1}^{S}\lambda_1\langle\tilde{A}^{i+1}, \tilde{Q}_A^{i+1}\rangle + \lambda_2\langle\tilde{B}^{i+1}, \tilde{Q}_B^{i+1}\rangle\\ &\quad + \frac{\mu_1}{2}(-\|\tilde{A}^1\|_F^2 + \sum_{i=1}^{S}\|\tilde{X}^{i+1} - \tilde{A}^{i+1}\|_F^2 + \|\tilde{X}^{i+1} - \tilde{A}^i\|_F^2 + \|A^S\|_F^2)\\ &\quad + \frac{\mu_2}{2}(-\|\tilde{B}^1\|_F^2 + \sum_{i=1}^{S}\|\tilde{X}^{i+1}P - \tilde{B}^{i+1}\|_F^2 + \|\tilde{X}^{i+1}P - \tilde{B}^i\|_F^2 + \|B^S\|_F^2)\ .\end{aligned}$$

The convexity of $\|.\|_1$ imply that the terms $< \tilde{A}^i, \tilde{Q_A}^i >$ and $< \tilde{B}^i, \tilde{Q_B}^i >$ are positives $(\forall i)$. μ_1, μ_2, λ_1 and λ_2 being non-negative, all terms of the above equations are non-negative and we have:

$$\begin{aligned}
&\frac{\mu_1}{2}(\|\tilde{D}_A^1\|_F^2 + \|\tilde{A}^1\|_F^2) + \frac{\mu_2}{2}(\|\tilde{D}_B^1\|_F^2 + \|\tilde{B}^1\|_F^2) \\
&\quad \geq \; 2\sum_{i=1}^{S} \|\Phi\tilde{X}^i\|_F^2 + \sum_{i=1}^{S} \lambda_1 \langle \tilde{A}^{i+1}, \tilde{Q}_A^{i+1} \rangle + \lambda_2 \langle \tilde{B}^{i+1}, \tilde{Q}_B^{i+1} \rangle \\
&\quad + \frac{\mu_1}{2}\|\tilde{X}^{i+1} - \tilde{A}^i\|_F^2 + \frac{\mu_2}{2}\|\tilde{X}^{i+1}P - \tilde{B}^i\|_F^2 \; .
\end{aligned}$$

From this last equation we can derive:

$$\sum_{i=1}^{\infty} \|\Phi\tilde{X}^i\|_F^2 < \infty \; , \tag{A.4}$$

$$\sum_{i=1}^{\infty} \langle \tilde{A}^{i+1}, \tilde{Q}_A^{i+1} \rangle \; < \infty, \; \sum_{i=1}^{\infty} \langle \tilde{B}^{i+1}, \tilde{Q}_B^{i+1} \rangle \; < \infty \; , \tag{A.5}$$

$$\sum_{i=1}^{\infty} \|\tilde{X}^{i+1} - \tilde{A}^i\|_F^2 \; < \infty, \; \sum_{i=1}^{\infty} \|\tilde{X}^{i+1}P - \tilde{B}^i\|_F^2 \; < \infty \; , \tag{A.6}$$

which leads to the convergence theorem enunciated in Section 4.2.
From Eq. (A.5) and the properties of the Bregman distance (cf. Eq. (3.6) in [7]):

$$\begin{aligned}
&\lim_{i\to\infty} \|A^i\|_1 - \|\hat{A}\|_1 - \langle A^i - \hat{A}, \; \hat{Q}_A \rangle = 0 \; , \\
&\lim_{i\to\infty} \|B^i\|_1 - \|\hat{B}\|_1 - \langle B^i - \hat{B}, \; \hat{Q}_B \rangle = 0 \; ,
\end{aligned}$$

which, combined with Eq. (A.6), gives:

$$\lim_{i\to\infty} \|X^i\|_1 - \|\hat{X}\|_1 - \langle X^i - \hat{X}, \; \hat{Q}_A \rangle = 0 \; , \tag{A.7}$$

$$\lim_{i\to\infty} \|X^iP\|_1 - \|\hat{X}P\|_1 - \langle X^i - \hat{X}, \; \hat{Q}_BP^T \rangle = 0 \; , \tag{A.8}$$

and finally, by taking λ_1 Eq. (A.7) + λ_2 Eq. (A.8) and Eq. (A.2), we have:

$$\begin{aligned}
&\lim_{i\to\infty} \|X^i\|_1 - \|\hat{X}\|_1 + \|X^iP\|_1 - \|\hat{X}P\|_1 \\
&\qquad - \langle X^i - \hat{X}, \; 2\Phi^T(Y - \Phi\hat{X}) \rangle = 0 \; .
\end{aligned} \tag{A.9}$$

In addition, $\|\Phi\tilde{X}^i\|_2^2 =< \nabla f(X^i) - \nabla f(\hat{X}), X^i - \hat{X} >$ for $f(X) = \|Y - \Phi X\|_2^2$. F is convex and from Eq. (A.4) and the same Bregman distance's property used before we can write:

$$\begin{aligned}
&\lim_{i\to\infty} \|Y - \Phi X^i\|_F^2 - \|Y - \Phi\hat{X}\|_F^2 \\
&\qquad - \langle X^i - \hat{X}, -2\Phi^T(Y - \Phi\hat{X}) \rangle \; ,
\end{aligned}$$

which provide with Eq. (A.9) the first result of the theorem:

$$\begin{aligned}\lim_{i\to\infty} &\|Y-\Phi X^i\|_F^2+\lambda_1\|X^i\|_1+\lambda_2\|X^iP\|_1\\ &=\|Y-\Phi\hat{X}\|_F^2+\lambda_1\|\hat{X}\|_1+\lambda_2\|\hat{X}P\|_1=0\ .\end{aligned}$$

The second part is obtained by noting that the fonction $g(X)=\|Y-\Phi X\|_2^2+\lambda_1\|X\|_1+\lambda_2\|XP\|_1$ is continuous, strictly convex and then has an unique minimizer. So, $\lim_{i\to\infty} g(X^i)=g(\hat{X})\Rightarrow\lim_{i\to\infty}X^i=\hat{X}$ (see Osher *et al.* [7]).

References

[1] Angelosante, D., Giannakis, G., Sidiropoulos, N., 2010. Multiple frequency-hopping signal estimation via sparse regression. In: Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on. IEEE, pp. 3502–3505.

[2] Bartels, R., Stewart, G., 1972. Solution of the matrix equation AX+ XB= C [F4]. Communications of the ACM 15 (9), 820–826.

[3] Beck, A., Teboulle, M., 2009. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. SIAM Journal on Imaging Sciences 2 (1), 183–202.

[4] Bertsekas, D., 1982. Constrained optimization and lagrange multiplier methods. Computer Science and Applied Mathematics, Boston: Academic Press 1.

[5] Blankertz, B., Müller, K.-R., Curio, G., Vaughan, T., Schalk, G., Wolpaw, J., Schlögl, A., Neuper, C., Pfurtscheller, G., Hinterberger, T., Schröder, M., Birbaumer, N., 2004. The BCI Competition 2003: progress and perspectives in detection and discrimination of EEG single trials. IEEE Trans. on Biomedical Engineering 51, 1044–1051.

[6] Bleakley, K., Vert, J.-P., 2011. The group fused Lasso for multiple change-point detection. arXiv preprint arXiv:1106.4199.

[7] Cai, J., Osher, S., Shen, Z., 2009. Split Bregman methods and frame based image restoration. Multiscale modeling and simulation 8 (2), 337.

[8] Candes, E. J., Eldar, Y. C., Needell, D., Randall, P., 2011. Compressed sensing with coherent and redundant dictionaries. Applied and Computational Harmonic Analysis 31 (1), 59–73.

[9] Chen, X., Lin, Q., Kim, S., Carbonell, J., Xing, E., 2012. Smoothing proximal gradient method for general structured sparse regression. The Annals of Applied Statistics 6 (2), 719–752.

[10] Combettes, P. L., Wajs, V. R., 2005. Signal recovery by proximal forward-backward splitting. Multiscale Modeling & Simulation 4 (4), 1168–1200.

[11] Cotter, S., Rao, B., Engan, K., Kreutz-Delgado, K., 2005. Sparse solutions to linear inverse problems with multiple measurement vectors. IEEE Trans. on Signal Processing 53 (7), 2477–2488.

[12] Danaher, P., Wang, P., Witten, D. M., 2014. The joint graphical Lasso for inverse covariance estimation across multiple classes. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 76 (2), 373–397.

[13] Darbon, J., Sigelle, M., 2005. A fast and exact algorithm for total variation minimization. In: Pattern recognition and image analysis. Vol. 3522 of Lecture Notes in Computer Science. pp. 351–359.

[14] Donoho, D., 2006. Compressed sensing. IEEE Trans. on Information Theory 52 (4), 1289–1306.

[15] Donoho, D., Elad, M., Temlyakov, V., 2006. Stable recovery of sparse overcomplete representations in the presence of noise. IEEE Trans. on Information Theory 52 (1), 6–18.

[16] Efron, B., Hastie, T., Johnstone, I., Tibshirani, R., 2004. Least angle regression. The Annals of statistics 32 (2), 407–499.

[17] Elad, M., Milanfar, P., Rubinstein, R., 2007. Analysis versus synthesis in signal priors. Inverse problems 23 (3), 947.

[18] Farwell, L. A., Donchin, E., 1988. Talking off the top of your head: toward a mental prosthesis utilizing event-related brain potentials. Electroencephalography and clinical Neurophysiology 70 (6), 510–523.

[19] Gholami, A., Siahkoohi, H., 2010. Regularization of linear and non-linear geophysical ill-posed problems with joint sparsity constraints. Geophysical Journal International 180 (2), 871–882.

[20] Goldstein, T., Osher, S., 2009. The split Bregman method for ℓ_1 regularized problems. SIAM Journal on Imaging Sciences 2 (2), 323–343.

[21] Gramfort, A., Papadopoulo, T., Olivi, E., Clerc, M., et al., 2010. OpenMEEG: opensource software for quasistatic bioelectromagnetics. Biomed. Eng. Online 9 (1), 45.

[22] Gramfort, A., Strohmeier, D., Haueisen, J., Hämäläinen, M. S., Kowalski, M., 2013. Time-frequency mixed-norm estimates: sparse M/EEG imaging with non-stationary source activations. NeuroImage 70, 410–422.

[23] Gribonval, R., Rauhut, H., Schnass, K., Vandergheynst, P., 2008. Atoms of all channels, unite! Average case analysis of multi-channel sparse recovery using greedy algorithms. Journal of Fourier analysis and Applications 14 (5), 655–687.

[24] Hoefling, H., 2010. A path algorithm for the fused Lasso signal approximator. Journal of Computational and Graphical Statistics 19 (4), 984–1006.

[25] Hoffmann, U., Vesin, J.-M., Ebrahimi, T., Diserens, K., 2008. An efficient P300-based brain–computer interface for disabled subjects. Journal of Neuroscience methods 167 (1), 115–125.

[26] Huang, J., Zhang, T., Metaxas, D., 2011. Learning with structured sparsity. Journal of Machine Learning Research 12, 3371–3412.

[27] Isaac, Y., Barthélemy, Q., Atif, J., Gouy-Pailler, C., Sebag, M., 2013. Multi-dimensional sparse structured signal approximation using split Bregman iterations. In: Acoustics Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on. IEEE, pp. 3826–3830.

[28] Isaac, Y., Barthélemy, Q., Gouy-Pailler, C., Atif, J., Sebag, M., 2015. Généralisation des micro-états EEG par apprentissage régularisé temporellement de dictionnaires topographiques. In: XXV Colloque GRETSI - Traitement du Signal et des Images.

[29] Jenatton, R., Audibert, J., Bach, F., 2011. Structured variable selection with sparsity-inducing norms. Journal of Machine Learning Research 12, 2777–2824.

[30] Jenatton, R., Mairal, J., Bach, F., Obozinski, G., 2010. Proximal methods for sparse hierarchical dictionary learning. In: Proceedings of the 27th International Conference on Machine Learning (ICML-10). pp. 487–494.

[31] Kim, S., Xing, E., 2009. Statistical estimation of correlated genome associations to a quantitative trait network. PLoS genetics 5 (8), e1000587.

[32] Kim, S.-J., Koh, K., Boyd, S., Gorinevsky, D., 2009. ℓ_1 trend filtering. Siam Review 51 (2), 339–360.

[33] Lee, T.-W., Lewicki, M., Girolami, M., Sejnowski, T., 1999. Blind source separation of more sources than mixtures using overcomplete representations. IEEE Signal Processing Letters 6 (4), 87–90.

[34] Liu, J., Yuan, L., Ye, J., 2010. An efficient algorithm for a class of fused Lasso problems. In: Proc. 16th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining. ACM, pp. 323–332.

[35] Mairal, J., Elad, M., Sapiro, G., 2008. Sparse representation for color image restoration. IEEE Trans. on Image Processing 17 (1), 53–69.

[36] Majumdar, A., Ward, R. K., 2012. Synthesis and analysis prior algorithms for joint-sparse recovery. In: Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on. IEEE, pp. 3421–3424.

[37] Nesterov, Y., 2005. Smooth minimization of non-smooth functions. Mathematical Programming 103 (1), 127–152.

[38] Pati, Y., Rezaiifar, R., Krishnaprasad, P., 1993. Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition. In: Signals, Systems and Computers, 1993. Conf. Record of The Twenty-Seventh Asilomar Conf. on. IEEE, pp. 40–44.

[39] Peyré, G., Fadili, J., 2011. Learning analysis sparsity priors. In: Sampta'11.

[40] Polich, J., 2007. Updating P300: an integrative theory of P3a and P3b. Clinical neurophysiology 118 (10), 2128–2148.

[41] Rakotomamonjy, A., 2011. Surveying and comparing simultaneous sparse approximation (or group-Lasso) algorithms. Signal Processing 91 (7), 1505–1526.

[42] Rubinstein, R., Faktor, T., Elad, M., 2012. K-SVD dictionary-learning for the analysis sparse model. In: Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on. IEEE, pp. 5405–5408.

[43] Rudin, L., Osher, S., Fatemi, E., 1992. Nonlinear total variation based noise removal algorithms. Physica D: Nonlinear Phenomena 60 (1-4), 259–268.

[44] Selesnick, I. W., Figueiredo, M. A., 2009. Signal restoration with overcomplete wavelet transforms: comparison of analysis and synthesis priors. In: SPIE Optical Engineering+ Applications. International Society for Optics and Photonics, pp. 74460D–74460D.

[45] Tibshirani, R., 1996. Regression shrinkage and selection via the Lasso. Journal of the Royal Statistical Society. Series B (Methodological) 58, 267–288.

[46] Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., Knight, K., 2005. Sparsity and smoothness via the fused Lasso. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 67 (1), 91–108.

[47] Tibshirani, R., Taylor, J., 2011. The solution path of the generalized Lasso. The Annals of Statistics 39 (3), 1335–1371.

[48] Tošić, I., Frossard, P., 2011. Dictionary learning. IEEE Signal Processing Magazine 28 (2), 27–38.

[49] Tropp, J., 2006. Algorithms for simultaneous sparse approximation. Part II: Convex relaxation. Signal Processing 86 (3), 589–602.

[50] Tropp, J., Gilbert, A., Strauss, M., 2006. Algorithms for simultaneous sparse approximation. Part I: Greedy pursuit. Signal Processing 86 (3), 572–588.

[51] Vaiter, S., Peyré, G., Dossal, C., Fadili, J., 2013. Robust sparse analysis regularization. IEEE Trans. on Information Theory 59 (4), 2001–2016.

[52] Valdés, P., Bosch, J., Grave, R., Hernandez, J., Riera, J., Pascual, R., Biscay, R., 1992. Frequency domain models of the EEG. Brain topography 4 (4), 309–319.

[53] Wahlberg, B., Boyd, S., Annergren, M., Wang, Y., 2012. An ADMM algorithm for a class of total variation regularized estimation problems. In: IFAC Symp. Syst. Ident. pp. 83–88.

[54] Wright, J., Yang, A., Ganesh, A., Sastry, S., Ma, Y., 2009. Robust face recognition via sparse representation. IEEE Trans. on Pattern Analysis and Machine Intelligence 31 (2), 210–227.

[55] Wu, C., Tai, X., 2010. Augmented Lagrangian method, dual methods, and split Bregman iteration for ROF, vectorial TV, and high order models. SIAM Journal on Imaging Sciences 3 (3), 300–339.

[56] Ye, G., Xie, X., 2011. Split Bregman method for large scale fused Lasso. Computational Statistics & Data Analysis 55 (4), 1552–1569.

[57] Yuan, M., Lin, Y., 2006. Model selection and estimation in regression with grouped variables. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 68 (1), 49–67.

[58] Zhou, J., Liu, J.and Narayan, V., Ye, J., 2012. Modeling disease progression via fused sparse group Lasso. In: Proc. 18th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining. ACM, pp. 1095–1103.